arXiv:1403.6303v1 [cs.LO] 25 Mar 2014

Verification for Timed Automata extended with Unbounded Discrete Data Structures

Karin Quaas*

Universität Leipzig, Germany

Abstract. We study decidability of verification problems for timed automata extended with unbounded discrete data structures. More detailed, we extend timed automata with a pushdown stack. In this way, we obtain a strong model that may for instance be used to model real-time programs with procedure calls. It is long known that the reachability problem for this model is decidable. The goal of this paper is to identify subclasses of timed pushdown automata for which the language inclusion problem and related problems are decidable.

1 Introduction

Timed automata were introduced by Alur and Dill [4], and have since then become a popular standard formalism to model real-time systems. An undeniable reason for the success of timed automata is the PSPACE decidability of the *language emptiness problem* [4]. A major drawback of timed automata is the undecidability [4] of the *language inclusion problem*:

Language Inclusion Problem
INPUT: Two timed automata $\mathcal{A}$ and $\mathcal{B}$.
QUESTION: Does $L(\mathcal{A}) \subseteq L(\mathcal{B})$ hold?

The undecidability of this problem prohibits the usage of automated verification algorithms for analysing timed automata, where $\mathcal{B}$ can be seen as the specification that is supposed to be satisfied by the system modelled by $\mathcal{A}$. It was, however, recently proved that the language inclusion problem is decidable (albeit with non-primitive recursive complexity) if $\mathcal{B}$ is restricted to have at most one clock [29]. Another milestone in the success story of timed automata is the decidability of the *model checking problem* for timed automata and Metric Temporal Logic (MTL, for short) over finite timed words [30].

Timed automata are capable to express many interesting time-related properties, and even with the restriction to a single clock, they allow to model a large class of systems, including, for example, the internet protocol TCP [28]. If we want to reason about real-time programs with procedure calls, or about the number of events occurring in computations of real-time systems, we have to

* The author is supported by *DFG*, project QU 316/1-1.

extend the model of timed automata with some unbounded discrete data structure. Bouajjani et al. [7] extended timed automata with discrete counters and a pushdown stack and proved that the satisfiability of reachability properties for several subclasses of this model is decidable. Nine years later, it was shown that the binary reachability relation for *timed pushdown systems* is decidable [14]. Decidability of the reachability problem was also proved for several classes of *timed counter systems* [8], mainly by simple extensions of the classical region-graph construction [4]. The language inclusion problem, however, was only considered in [18] for the class of timed pushdown systems. In [18] it is stated that the language inclusion problem is decidable if $\mathcal{A}$ is a timed pushdown automaton, and $\mathcal{B}$ is a single-clock timed automaton. The proof is based on an extension of the proof for the decidability of language inclusion for timed automata in [29]. Unfortunately, and as is well known, the proof in [18] is not correct.

In this paper, we prove that different to what is claimed in [18], the language inclusion problem for the case that $\mathcal{A}$ is a pushdown timed automaton and $\mathcal{B}$ is a single-clock timed automaton is undecidable. This is even the case if $\mathcal{A}$ is a deterministic instance of a very restricted subclass of timed pushdown automata called *timed visibly one-counter nets*. On the other hand, we prove that the language inclusion problem is decidable if $\mathcal{A}$ is a timed automaton and $\mathcal{B}$ is a timed automaton extended with a finite set of counters that can be incremented and decremented, and which we call *timed counter nets*. As a special case, we obtain the decidability of the *universality problem* for timed counter nets:

Universality Problem
INPUT: A timed automaton $\mathcal{A}$ with input alphabet Σ.
QUESTION: Does $L(\mathcal{A})$ recognize the set of all timed words over Σ?

Finally, we give the precise decidability border for the universality problem by proving that the universality problem is undecidable for the class of *timed visibly one-counter automata*. All our results apply to extensions of timed automata over *finite* timed words.

2 Extensions of Timed Automata with Discrete Data Structure

We use $\mathbb{Z}$, $\mathbb{N}$ and $\mathbb{R}_{\geq 0}$ to denote the integers, the non-negative integers and the non-negative reals, respectively.

We use Σ to denote a finite alphabet. A *timed word* over Σ is a non-empty finite sequence $(a_1, t_1) \dots (a_k, t_n) \in (\Sigma \times \mathbb{R}_{\geq 0})^+$ such that the sequence $t_1, \dots, t_n$ of timestamps is non-decreasing. We say that a timed word is *strictly monotonic* if $t_{i-1} < t_i$ for every $i \in \{2, \dots, n\}$. We use $T\Sigma^+$ to denote the set of finite timed words over Σ. A set $L \subseteq T\Sigma^+$ is called a *timed language*. We use $|w|$ to denote the *length* of a timed word w.

Let $\mathcal{X}$ be a finite set of *clock variables* ranging over $\mathbb{R}_{\geq 0}$. We define *clock constraints* ϕ over $\mathcal{X}$ to be conjunctions of formulas of the form $x \sim c$, where

$x \in \mathcal{X}$, $c \in \mathbb{N}$, and $\sim \in \{<, \leq, =, \geq, >\}$. We use $\Phi(\mathcal{X})$ to denote the set of all clock constraints over $\mathcal{X}$. A *clock valuation* is a mapping from $\mathcal{X}$ to $\mathbb{R}_{\geq 0}$. A clock valuation ν satisfies a clock constraint ϕ, written $\nu \models \phi$, if ϕ evaluates to true according to the values given by ν. For $\delta \in \mathbb{R}_{\geq 0}$ and $\lambda \subseteq \mathcal{X}$, we define $\nu + \delta$ to be $(\nu + \delta)(x) = \nu(x) + \delta$ for each $x \in \mathcal{X}$, and we define $\nu[\mathcal{X} := 0]$ by $(\nu[\mathcal{X} := 0])(x) = 0$ if $x \in \mathcal{X}$, and $(\nu[\mathcal{X} := 0])(x) = \nu(x)$ otherwise.

Let Γ be a finite stack alphabet. We use Γ^* to denote the set of finite words over Γ, including the empty word ε. We define a finite set Op of *stack operations* by $\mathsf{Op} := \{pop(a), push(a) \mid a \in \Gamma\} \cup \{noop, empty?\}$.

A *timed pushdown automaton* is a tuple $\mathcal{A} = (\Sigma, \Gamma, \mathcal{L}, \mathcal{L}_0, \mathcal{L}_f, \mathcal{X}, E)$, where

- $\mathcal{L}$ is a finite set of *locations*,
- $\mathcal{L}_0 \subseteq \mathcal{L}$ is the set of initial locations,
- $\mathcal{L}_f \subseteq \mathcal{L}$ is the set of accepting locations,
- $E \subseteq \mathcal{L} \times \Sigma \times \Phi(\mathcal{X}) \times \mathsf{Op} \times 2^{\mathcal{X}} \times \mathcal{L}$ is a finite set of edges.

Each edge $(l, a, \phi, \mathsf{op}, \lambda, l')$ represents a discrete transition from l to l' on the input symbol a under certain conditions: The clock constraint ϕ and the set λ specify the constraints on the current values of the clocks and the clocks to be reset, respectively, and op specifies the operations on the stack.

A *state* of $\mathcal{A}$ is a triple (l, ν, u), where $l \in \mathcal{L}$ is the current location, the clock valuation ν represents the current values of the clock, and $u \in \Gamma^*$ represents the current stack content, where the top-most symbol of the stack is the left-most symbol in the word u, and the empty word ε represents the empty stack. We use $\mathcal{G}^{\mathcal{A}}$ to denote the set of all states of $\mathcal{A}$. A timed pushdown automaton $\mathcal{A}$ induces a transition relation $\Rightarrow_{\mathcal{A}}$ on $(\mathcal{G}^{\mathcal{A}} \times \mathbb{R}_{\geq 0} \times \Sigma \times \mathcal{G}^{\mathcal{A}})$ as follows: $\langle (l, \nu, u), \delta, a, (l', \nu', u') \rangle \in \Rightarrow_{\mathcal{A}}$, if, and only if, there exists some edge $(l, a, \phi, \mathsf{op}, \lambda, l') \in E$ such that $(\nu + \delta) \models \phi$, $\nu' = (\nu + \delta)[\lambda := 0]$, and (i) if $\mathsf{op} = pop(a)$ for some $a \in \Gamma$, then a is the top-most stack symbol, *i.e.*, there exists $v \in \Gamma^*$ such that $u = a \cdot u'$; (ii) if $\mathsf{op} = push(a)$ for some $a \in \Gamma$, then a is pushed on the stack, *i.e.*, $u' = a \cdot u$; (iii) if $\mathsf{op} = empty?$, then the stack is empty and is is not changed by the operation, *i.e.*, $u = \varepsilon$, and $u' = u$; (iv) if $\mathsf{op} = noop$, then the stack is not changed, *i.e.*, $u' = u$. A *run* of $\mathcal{A}$ is a finite sequence $\prod_{1 \leq i \leq n} \langle (l_{i-1}, \nu_{i-1}, u_{i-1}), \delta_i, a_i, (l_i, \nu_i, u_i) \rangle$ such that $\langle (l_{i-1}, \nu_{i-1}, u_{i-1}), \delta_i, a_i, (l_i, \nu_i, u_i) \rangle \in \Rightarrow_{\mathcal{A}}$ for every $i \in \{1, \ldots, n\}$. A run is called *successful* if $l_0 \in \mathcal{L}_0$, $\nu_0(x) = 0$ for every $x \in \mathcal{X}$, $u_0 = \varepsilon$, and $l_n \in \mathcal{L}_f$. With a run we associate the timed word $(a_1, \delta_1)(a_2, \delta_1 + \delta_2) \ldots (a_n, \Sigma_{1 \leq i \leq n} \delta_i)$. The language accepted by a timed automaton, denoted by $L(\mathcal{A})$, is defined to be the set of timed words $w \in T\Sigma^+$ for which there exists a successful run of $\mathcal{A}$ that w is associated with.

Next we define some subclasses of timed pushdown automata; see also Fig. 1 for a graphical overview. We start with timed extensions of *one-counter automata* [15,26] and *one-counter nets* [23,1]. A *timed one-counter automaton* is a timed pushdown automaton where the stack alphabet is a singleton. By writing *push* and *pop* we mean that we increment and decrement the counter, respectively, whereas *empty*? corresponds to a zero test. A *timed one-counter net* is a timed one-counter automaton without zero tests, *i.e.*, the *empty*? operation

is not allowed. We remark that for both classes, the execution of an edge of the form $(l, a, \phi, pop, \lambda, l')$ is *blocked* if the stack is empty. Next, we consider the timed extension of an interesting subclass of pushdown automata called *visibly pushdown automata* [5]. A *timed visibly pushdown automaton* is a timed pushdown automaton for which the input alphabet Σ can be partitioned into three disjoint sets $\Sigma = \Sigma_{\mathsf{int}} \cup \Sigma_{\mathsf{call}} \cup \Sigma_{\mathsf{ret}}$ of *internal*, *call*, and *return* input symbols, respectively, and such that for every edge $(l, a, \phi, \mathsf{op}, \lambda, l')$ the following conditions are satisfied:

- $a \in \Sigma_{\mathsf{int}}$ if, and only if, $\mathsf{op} = noop$,
- $a \in \Sigma_{\mathsf{call}}$, if, and only if, $\mathsf{op} = push(b)$ for some $b \in \Gamma$,
- $a \in \Sigma_{\mathsf{ret}}$ if, and only if, $\mathsf{op} = empty?$ or $\mathsf{op} = pop(b)$ for some $b \in \Gamma$.

A *timed visibly one-counter automaton* (timed visibly one-counter net, respectively) is a timed one-counter automaton (timed one-counter net, respectively) that is also a timed visibly pushdown automaton. We say that a timed visibly one-counter net with no clock is *deterministic* if for all $e = (l, a, \mathtt{true}, \mathsf{op}', \emptyset, l'), e' = (l, a, \mathtt{true}, \mathsf{op}'', \emptyset, l'') \in E$ with $e \neq e'$ we have either $\mathsf{op}' = pop$ and $\mathsf{op}'' = empty?$, or $\mathsf{op}' = empty?$ and $\mathsf{op}'' = pop$.

Finally, we define the class of *timed counter nets*, which generalizes timed one-counter nets, but is not a subclass of timed pushdown automata. A timed counter net of dimension n is a tuple $\mathcal{A} = (\Sigma, n, \mathcal{L}, \mathcal{L}_0, \mathcal{L}_f, \mathcal{X}, E)$, where $\mathcal{L}, \mathcal{L}_0, \mathcal{L}_f$ are the sets of locations, initial locations and accepting locations, respectively, and $E \subseteq \mathcal{L} \times \Sigma \times \Phi(\mathcal{X}) \times \{0, 1, -1\}^n \times 2^{\mathcal{X}} \times \mathcal{L}$ is a finite set of edges. A state of a timed counter net is a triple $(l, \nu, \boldsymbol{x})$, where $l \in \mathcal{L}$, ν is a clock valuation, and $\boldsymbol{x} \in \mathbb{N}^n$ is a vector representing the current value of the counters. We define $\langle (l, \nu, \boldsymbol{x}), \delta, a, (l', \nu', \boldsymbol{x'}) \rangle \in \Rightarrow_{\mathcal{A}}$ if, and only if, there exists some edge $(l, a, \phi, \boldsymbol{c}, \lambda, l') \in E$ such that $(\nu + \delta) \models \phi$, $\nu' = (\nu + \delta)[\lambda := 0]$, and $\boldsymbol{x'} = \boldsymbol{x} + \boldsymbol{c}$, where vector addition is defined pointwise. Note that, similar to pop operations on an empty stack, transitions which result in the negative value of one of the counters are *blocked*. The notions of *runs*, *successful runs*, *associated timed words* and *the language accepted by* $\mathcal{A}$, are defined analogously to the corresponding definitions for timed pushdown automata.

3 Main Results

In this section, we present the main results of the paper. We are interested in the language inclusion problem $L(\mathcal{A}) \subseteq L(\mathcal{B})$, where $\mathcal{A}$ and $\mathcal{B}$ are extensions of timed automata with discrete data structure. Recall that according to standard notation in the field of verification, in this problem formulation $\mathcal{B}$ is seen as the *specification*, and $\mathcal{A}$ is the system that should satisfy this specification, *i.e.*, $\mathcal{A}$ should be a *model* of $\mathcal{B}$. As a special case of this problem, we also consider the universality problem, *i.e.*, the question whether $L(\mathcal{A}) = T\Sigma^+$ for a given automaton $\mathcal{A}$. In general, the two problems are undecidable for timed pushdown automata. This follows on the one hand side from the undecidability of the universality problem for timed automata [4], and on the other hand side from

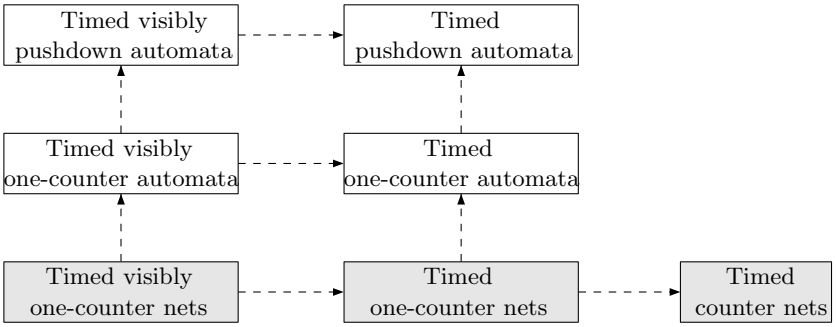

Fig. 1. Extensions of timed automata with discrete data structures. The subclass relation is represented by dashed arrows. The language emptiness problem is decidable for all classes. The classes in grey boxes have a decidable universality problem, the classes in white boxes bave an undecidable universality problem.

the undecidability of the universality problem for pushdown automata. In fact, it has been long known that the universality problem is undecidable already for non-deterministic one-counter automata [21,24].

However, there are interesting decidability results for subclasses of timed pushdown automata: The language inclusion problem is decidable if $\mathcal{A}$ is a timed automaton, and $\mathcal{B}$ is a timed automaton with at most one clock [29]. As a special case, the universality problem for timed automata is decidable if only one clock is used. The language inclusion problem is also decidable if $\mathcal{A}$ is a one-counter net and $\mathcal{B}$ is a finite automaton, and if $\mathcal{A}$ is a finite automaton and $\mathcal{B}$ is a one-counter net [25]. This also implies the decidability of the universality problem for one-counter nets. Further we know that the universality and language inclusion problems are decidable if $\mathcal{A}$ and $\mathcal{B}$ are visibly pushdown automata [6].

Hence it is interesting to consider the two problems for the corresponding subclasses of timed pushdown automata. It turns out that the decidability status changes depending from whether the *model* uses a stack (or, more detailed: a counter) or not. As a first main result, we have:

Theorem 1. *The language inclusion problem is undecidable if $\mathcal{A}$ is a timed visibly one-counter net and $\mathcal{B}$ is a timed automaton, even if $\mathcal{A}$ is deterministic and has no clocks, and $\mathcal{B}$ uses at most one clock.*

We remark that this result corrects a claim concerning the decidability of the language inclusion problem if $\mathcal{A}$ is a timed pushdown automaton and $\mathcal{B}$ is a single-clock timed automaton, stated in Theorem 2 in [18]. In contrast to this, we have decidability for the following classes:

Theorem 2. *The language inclusion problem is decidable with non-primitive recursive complexity if $\mathcal{A}$ is a timed automaton and $\mathcal{B}$ is a single-clock timed counter net.*

As a special case of this result, we obtain:

Corollary 1. *The universality problem for single-clock timed counter nets is decidable.*

The next two sections are devoted to the proofs of Theorems 1 and 2. We will also give some interesting consequences of these results respectively of their proofs. Amongst others, we prove for instance the undecidability of model checking problem for timed visibly one-counter nets and MTL over finite timed words. After this, in Sect. 5, we will prove the following theorem:

Theorem 3. *The universality problem for single-clock timed visibly one-counter automata is undecidable.*

This is in contrast to the decidability of the universality problem for the two underlying models of single-clock timed automata [29] and visibly one-counter automata, which is a subclass of visibly pushdown automata [6]. We also want to point out that this result is stronger than a previous result on the undecidability of the universality problem for single-clock timed visibly pushdown automata (Theorem 3 in [18]), and our proof closes a gap in the proof for Theorem 3 in [18]. Further, we can infer from Corollary 1 and Theorem 3 the exact decidability border for the universality problem of timed pushdown automata, which lies between timed visibly one-counter nets and timed visibly one-counter automata.

4 Undecidability Results

In this section, we prove Th. 1. The proof is a reduction of an undecidable problem for *channel machines*.

4.1 Channel Machines

Let A be a finite alphabet. We use ε to denote the *empty word* over A. Given two finite words $x, y \in A^*$, we use $x \cdot y$ to denote the *concatenation* of x any y. We define the order $\leq$ over the set of finite words over A by $x_1 x_2 \dots x_m \leq y_1 y_2 \dots y_n$ if there exists a strictly increasing function $f : \{1, \dots, m\} \to \{1, \dots, n\}$ such that $x_i = y_{f(i)}$ for every $i \in \{1, \dots, m\}$.

A *channel machine* consists of a finite-state automaton acting on an unbounded fifo channel. Formally, a channel machine is a tuple $\mathcal{C} = (S, s_I, M, \Delta)$, where

- S is a finite set of *control states*,
- $s_I \in S$ is the initial control state,
- M is a finite set of *messages*,
- $\Delta \subseteq S \times L \times S$ is the transition relation over the label set $L = \{!m, ?m \mid m \in M\} \cup \{\varepsilon\}$.

Here, $!m$ corresponds to a *send* operation, $?m$ corresponds to a *read* operation, and ε is a test which returns `true` if and only if the channel is empty. Without loss of generality, we assume that s_I does not have any incoming transitions,

i.e., $(s, l, s') \in \Delta$ implies $s' \neq s_I$. Further, we assume that $(s_I, l, s') \in \Delta$ implies $l = \varepsilon$. A *configuration* of $\mathcal{C}$ is a pair (s, x), where $s \in S$ is the control state and $x \in M^*$ represents the contents of the channel. We use $\mathcal{H}^{\mathcal{C}}$ to denote the set of all configurations of $\mathcal{C}$. The rules in Δ induce a transition relation $\rightarrow_{\mathcal{C}}$ on $(\mathcal{H}^{\mathcal{C}} \times L \times \mathcal{H}^{\mathcal{C}})$ as follows:

- $\langle (s, x), !m, (s', x') \rangle \in \rightarrow_{\mathcal{C}}$ if, and only if, there exists some transition $(s, !m, s') \in \Delta$ and $x' = x \cdot m$, *i.e.*, m is added to the tail of the channel.
- $\langle (s, x), ?m, (s', x') \rangle \in \rightarrow_{\mathcal{C}}$ if, and only if, there exists some transition $(s, ?m, s') \in \Delta$ and $x = m \cdot x'$, *i.e.*, m is the head of the current channel content.
- $\langle (s, x), \varepsilon, (s', x') \rangle \in \rightarrow_{\mathcal{C}}$ if, and only if, there exists some transition $(s, \varepsilon, s') \in \Delta$ and $x = \varepsilon$, *i.e.*, the channel is empty, and $x' = x$.

Next, we define a second transition relation $\rightsquigarrow_{\mathcal{C}}$ on $(\mathcal{H}^{\mathcal{C}} \times L \times \mathcal{H}^{\mathcal{C}})$. The relation $\rightsquigarrow_{\mathcal{C}}$ is a superset of $\rightarrow_{\mathcal{C}}$. It contains some additional transitions which result from *insertion errors*. We define $\langle (s, x_1), l, (s, x_1') \rangle \in \rightsquigarrow_{\mathcal{C}}$, if, and only if, there exist $x, x' \in M^*$ such that $\langle (s, x), l, (s', x') \rangle \in \rightarrow_{\mathcal{C}}$, $x_1 \leq x$, and $x' \leq x_1'$. A computation of $\mathcal{C}$ is a finite sequence $\prod_{1 \leq i \leq k} \langle (s_{i-1}, x_{i-1}), l_i, (s_i, x_i) \rangle$ such that $\langle (s_{i-1}, x_{i-1}), l_i, (s_i, x_i) \rangle \in \rightsquigarrow_{\mathcal{C}}$ for every $i \in \{1, \ldots, k\}$. We say that a computation is *error-free* if for all $i \in \{1, \ldots, k\}$ we have $\langle (s_{i-1}, x_{i-1}), l_i, (s_i, x_i) \rangle \in \rightarrow_{\mathcal{C}}$. Otherwise, we say that the computation is *faulty*.

Control State Reachability Problem for Channel Machines
INPUT: A channel machine $\mathcal{C}$ with control states S, control state $s_F \in S$.
QUESTION: Is there an error-free computation of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$?

The control state reachability problem is undecidable for channel machines, because channel machines are Turing-powerful [12,2]. The proof of Th. 1 is a reduction of the control state reachability problem for channel machines. The idea of the reduction is as follows: Given a channel machine $\mathcal{C}$, we define a timed language $L(\mathcal{C})$ consisting of all timed words that encode potentially faulty computations of $\mathcal{C}$ that start in (s_I, ε) and end in s_F. Then we define a timed visibly one-counter net $\mathcal{A}$ such that $L(\mathcal{A}) \cap L(\mathcal{C})$ contains all encodings of error-free such computations. In other words, we use $\mathcal{A}$ to exclude the encodings of faulty computations from $L(\mathcal{C})$. Finally, we define a single-clock timed automaton $\mathcal{B}$ that recognizes the complement of $L(\mathcal{C})$.

4.2 Encoding Faulty Computations

For the remainder of Section 4, let $\mathcal{C} = (S, s_I, M, \Delta)$ be a channel machine and let $s_F \in S$. Define $\Sigma_{\mathsf{int}} := S \cup M \cup L \backslash \{s_I\} \cup \{\#\}$, $\Sigma_{\mathsf{call}} := \{s_I, +\}$, and $\Sigma_{\mathsf{ret}} := \{-, \star\}$, where $+, -, \#$ and $\star$ are fresh symbols that do not occur in $S \cup M \cup L$. We define a timed language $L(\mathcal{C})$ over $\Sigma = \Sigma_{\mathsf{int}} \cup \Sigma_{\mathsf{call}} \cup \Sigma_{\mathsf{ret}}$ that consists of all timed words that encode (potentially faulty) computations of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$. The definition of $L(\mathcal{C})$ follows the ideas

presented in [30] to show the non-primitive recursive complexity of the model checking problem for MTL over finite timed words.

Let $\gamma = \prod_{1 \leq i \leq k} \langle (s_{i-1}, x_{i-1}), l_i, (s_i, x_i) \rangle$ be a computation of $\mathcal{C}$ with $s_0 = s_I$, $x_0 = \varepsilon$, and $s_k = s_F$. For each $i \in \{0, \ldots, k\}$, the configuration (s_i, x_i) is encoded by a timed word of duration one. This timed word starts with the symbol s_i at some time t_i. If the content of the channel x_i is of the form $m_1 m_2 \ldots m_j$, then s_i is followed by the symbols $m_1, m_2, \ldots, m_j$ in this order. The timestamps of these symbols must be in the interval $(t_i, t_i + 1)$. Note that due to the denseness of the time domain, one can indeed store the channel content in one time unit without any upper bound on j. For encoding the computation, we glue together the encodings of the single configurations as follows: Every control state symbol s_{i-1} is followed by l_i after exactly one time unit, and by s_i after exactly two time units. For every message symbol m between s_{i-1} and l_i, there is a copy of m after two time units, unless it is removed from the channel by a read operation. There are no symbols between l_i and s_i. Note that we *do not* require that for every message symbol m between s_{i-1} and l_i there is a copy of m two time units before. It is the absence of exactly this condition that causes $L(\mathcal{C})$ to contain encodings of computations that are potentially faulty.

For our reduction to work, we change the idea from [30] in some details. As mentioned above, we will later define a timed visibly one-counter net to exclude faulty computations from $L(\mathcal{C})$. The idea is to let the automaton *guess* the maximum number n of messages occurring between control state symbols and the following label symbol. While reading the encoding of the first configuration of a computation, it increments the counter n times. The automaton does not do any operations on the counter until it reads the encoding of the last configuration, where it decrements the counter whenever it reads message symbols occurring between the control state and the label symbol. Since it can decrement the counter at most n times, it can only accept encodings of computations without any insertion errors.

We define a timed language $L(\mathcal{C}, n)$ for every $n \in \mathbb{N}$. Since *push*, *pop* and *noop* require symbols from the call, internal, and return input alphabet, respectively, we use three different extra symbols $+, \#$, and $-$ to represent the non-deterministic choice of the automaton. In the encoding of the initial configuration, we use n occurrences of the call symbol $+$ as placeholder for message symbols. In the encoding of all following configurations except for the last one we use the internal symbol $\#$ as placeholder. In the last configuration, we use the return symbol $-$ as placeholder. Every time a send operation of message m occurs, the first free placeholder symbol in the encoding of the current configuration is replaced by m in the encoding of the next configuration. Every time a read operation of message m occurs, the first message symbol in the encoding of the current configuration (which should be m) is replaced by a new placeholder symbol in the encoding of the next configuration. Special cases due to insertion errors or the number of messages between some control state and label symbol are treated as defined in what follows.

Let $n \in \mathbb{N}$. The timed language $L(\mathcal{C}, n)$ consists of all timed words w over Σ that satisfy the following conditions:

1. w must be strictly monotonic.
2. In w, for every control state symbol s different from s_F there exists a label symbol l after one time unit, and a control state symbol s' after two time units, provided that $(s, l, s') \in \Delta$. For the symbol s_F there exists the symbol $\star$ after one time unit. Control state symbols, label symbols and the symbol $\star$ must not occur anywhere else in w.
3. Message symbols may occur in w between a control state symbol different from s_I and s_F and a label symbol. They must not occur anywhere else in w.
4. The placeholder symbol $+$ may only occur between s_I and the next label symbol.
5. The placeholder symbol $\#$ may only occur between control state symbols different from s_I and s_F and the next label symbol.
6. The placeholder symbol $-$ may only occur between s_F and the next $\star$ symbol.
7. Between a control state and a label symbol, hash symbols $\#$ may only occur *after* message symbols.
8. The (untimed) prefix of w must be of the form $s_I +^n ls$ for some $l \in L, s \in S$.
9. For every $+$, there is a $\#$ after two time units[1].
10. The (untimed) suffix of w must be of the form $s_F -^m \star$ for some $m \in \mathbb{N}$.

Further, for every infix of w of the form

$$(s, \delta)(\sigma_1, \delta_1)(\sigma_2, \delta_2) \dots (\sigma_k, \delta_k)(l, \delta + 1)(s', \delta + 2)(\sigma'_1, \delta'_1) \dots (\sigma'_{k'}, \delta'_{k'})(l', \delta + 3)$$

with $s, s' \in S \backslash \{s_I, s_F\}$, $l, l' \in L$, $\delta \in \mathbb{R}_{\geq 0}$, $\delta < \delta_1 < \dots < \delta_k < \delta + 1$, $\delta + 2 < \delta'_1 < \dots < \delta'_{k'} < \delta + 3$, there exists a strictly increasing function $f : \{1, \dots, k\} \to \{1, \dots, k'\}$ such that the following conditions are satisfied:

11. If $l = \varepsilon$, then $\sigma_i = \#$, $\sigma'_{f(i)} = \sigma_i$ and $\delta'_{f(i)} = \delta_i + 2$ for every $i \in \{1, \dots, k\}$. (It is tested whether the channel is empty, *i.e.*, whether there exists only placeholder symbols between s and l, and there exists a copy of every symbol after two time units.)
12. If $l = !m$ for some $m \in M$, then one of the following conditions holds:
    (a) If there exists $i \in \{1, \dots, k\}$ such that $\sigma_i = \#$, then $\sigma'_{f(j)} = m$ for $j = \min\{i \in \{1, \dots, k\} \mid \sigma_i = \#\}$, $\sigma'_{f(i)} = \sigma_i$ for all $i \in \{1, \dots, k\} \backslash \{j\}$, and $\delta'_{f(i)} = \delta_i + 2$ for every $i \in \{1, \dots, k\}$. (The first placeholder symbol between s and l, if there is some, is replaced by m two time units later. For all other symbols there is a copy two time units later.)
    (b) If $\sigma_i \neq \#$ for all $i \in \{1, \dots, k\}$, then $k' > k$, $\sigma'_{f(i)} = \sigma_i$, $\delta'_{f(i)} = \delta_i + 2$ for all $i \in \{1, \dots, k\}$, and there exists $j > f(k)$ such that $\sigma'_j = \#$. (If there exists no placeholder symbol between s and l, then there exists a copy for each symbol after two time units, and there exists some additional placeholder symbol after the last copied symbol. Note that this case corresponds to n being chosen too small.)

[1] Recall that every transition $(s_I, l, s) \in \Delta$ satisfies $l = \varepsilon$.

13. If $l = ?m$ for some $m \in M$, then one of the following conditions holds:
    (a) If $\sigma_1 = m$, then $\sigma'_{f(i)} = \sigma_{i+1}$, $\delta'_{f(i)} = \delta_i + 2$ for all $i \in \{1, \ldots, k-1\}$, and $\sigma'_{f(k)} = \#$, $\delta'_{f(k)} = \delta_k + 2$. (If the first symbol between s and l is m, then there exists a new placeholder symbol after the copy of all other symbols, replacing the first m symbol. Timestamps are shifted to the right.)
    (b) If $\sigma_1 \neq m$, then $k' > k$, $\sigma'_{f(i)} = \sigma_i$, $\delta'_{f(i)} = \delta_i + 2$ for all $i \in \{1, \ldots, k\}$, and there exists some $j > f(k)$ such that $\sigma'_j = \#$. (If the first symbol is not m, then, besides a copy of all symbols between s and l, there exists a new placeholder symbol between s' and l'. Note that this case corresponds to an insertion error.)

For an infix like above but with $s' = s_F$ and $l' = \star$, we add conditions 11', 12' and 13' that differ from 11, 12, and 13, respectively, in that all message or placeholder symbols to be copied or added to the encoding of the next configuration are replaced by the placeholder symbol $-$. This finishes the definition of $L(\mathcal{C}, n)$.

Let $w \in L(\mathcal{C}, n)$ for some $n \in \mathbb{N}$. We use $\max(w)$ to denote the maximum number of symbols in $M \cup \{\#, +, -\}$ that occur in w between a control state symbol and a symbol in $L \cup \{\star\}$. Let $\gamma = \prod_{1 \leq i \leq k} \langle (s_{i-1}, x_{i-1}), l_i, (s_i, x_i) \rangle$ be a computation of $\mathcal{C}$. We use $\max(\gamma)$ to denote the maximum length of the channel content occurring in γ, formally: $\max(\gamma) := \max\{|x_i| \mid 0 \leq x_i \leq k\}$.

Lemma 1. *For each error-free computation γ of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$, there exists some timed word $w \in L(\mathcal{C}, \max(\gamma))$ such that $\max(w) = \max(\gamma)$.*

Proof. Let γ be an error-free computation of $\mathcal{C}$ from $= (s_I, \varepsilon)$ to (s_F, x) for some $x \in M^*$. Put $n = \max(\gamma)$. Clearly, there exists some $w \in L(\mathcal{C}, n)$, but we have to prove that adding message or placeholder symbols is not necessary, *i.e.*, there exists $w \in L(\mathcal{C}, n)$ such that $\max(w) = n$. By the choice of n and the fact that γ is error-free, the conditions of the only cases where adding symbols is mandatory (namely 12(b) and 13(b) and their primed variants) do not apply. □

Lemma 2. *For each $n \in \mathbb{N}$ and $w \in L(\mathcal{C}, n)$ with $\max(w) = n$, there exists some error-free computation γ of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$ with $\max(\gamma) \leq n$.*

Proof. Let $n \in \mathbb{N}$ and let $w \in L(\mathcal{C}, n)$ with $\max(w) = n$. This implies that the conditions of 12(b), 13(b) and their primed variants are never satisfied. Hence there exists some error-free computation γ of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$ with $\max(\gamma) \leq n$. □

4.3 Excluding Faulty Computations

We define a timed visibly one-counter net $\mathcal{A}$ over Σ such that for every $n \in \mathbb{N}$ the intersection $L(\mathcal{A}) \cap L(\mathcal{C}, n)$ consists of all timed words that encode *error-free* computations of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$. The timed visibly

one-counter net $\mathcal{A}$ is shown in Fig. 4. It non-deterministically guesses a number $n \in \mathbb{N}$ of symbols $+$ and increments the counter each time it reads the symbol $+$. When $\mathcal{A}$ leaves l_1, the value of the counter is $n+1$. After that, the counter value is not changed until the state symbol s_F is read. Then, while reading symbols in $\{-, \star\}$, the counter value is decremented. Note that $\mathcal{A}$ can reach the final location l_4 only if the number of the occurrences of symbol $-$ between s_F and $\star$ is *at most* n. Note that $\mathcal{A}$ does not use any clock, and it is deterministic.

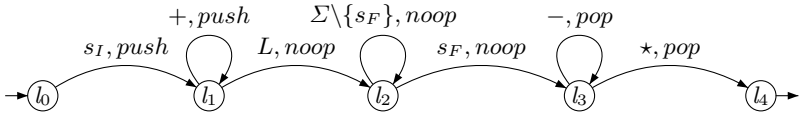

Fig. 2. The deterministic timed visibly one-counter net $\mathcal{A}$ for excluding insertion errors.

Lemma 3. *$\mathcal{C}$ has an error-free computation from (s_I, ε) to (s_F, x) for some $x \in M^*$, if, and only if, there exists $n \in \mathbb{N}$ such that $L(\mathcal{C}, n) \cap L(\mathcal{A}) \neq \emptyset$.*

Proof. For the direction from left to right, let $\gamma = \prod_{1 \leq i \leq k} \langle (s_{i-1}, x_{i-1}), l_i, (s_i, x_i) \rangle$ be an error-free computation of $\mathcal{C}$ such that $s_0 = s_I$, $x_0 = \varepsilon$, and $s_k = s_F$. Define $n = \max(\gamma)$. By Lemma 1, there exists $w \in L(\mathcal{C}, n)$ such that $\max(w) = n$. This implies $w \in L(\mathcal{A})$. Hence $L(\mathcal{C}, n) \cap L(\mathcal{A}) \neq \emptyset$.

For the direction from right to left, let $w \in L(\mathcal{C}, n) \cap L(\mathcal{A})$ for some $n \in \mathbb{N}$. By definition of $\mathcal{A}$, the value of the counter after reading the (untimed) prefix $s_I +^n$, equals $n+1$. By definition of $L(\mathcal{C}, n)$ (conditions 11, 12, 13, 11', 12', 13'), the (untimed) suffix of w is of the form $s_F -^m \star$ for some $m \geq n$. However, by definition of $\mathcal{A}$, $m \leq n$, because otherwise the edge to l_4 cannot be taken. Hence $m = n$ and $\max(w) = n$. By Lemma 2, there exists some error-free computation γ of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$ with $\max(\gamma) \leq n$. □

We finally define $L(\mathcal{C}) = \bigcup_{n \in \mathbb{N}} L(\mathcal{C}, n)$.

Corollary 2. *There exists some error-free computation of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$ if, and only if, $L(\mathcal{A}) \cap L(\mathcal{C}) \neq \emptyset$.*

4.4 The Reduction

Finally, we define a single-clock timed automaton $\mathcal{B}$ such that $L(\mathcal{B}) = T\Sigma^+ \setminus L(\mathcal{C})$. The construction of $\mathcal{B}$ follows the same ideas as, *eg.*, in [3]: $\mathcal{B}$ is the union of several single-clock timed automata, each of them violating one of the conditions of the definition of $L(\mathcal{C})$. For instance, the timed automaton in Fig. 2 recognizes the set of timed words over Σ violating condition 2.

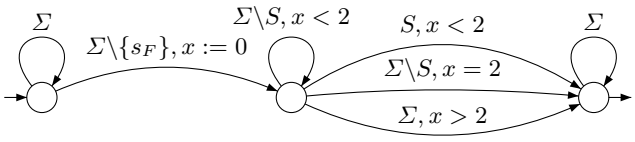

Fig. 3. A timed automaton violating condition 2 of $L(\mathcal{C})$

Proof of Theorem 1 Let $\mathcal{C} = (S, s_0, M, \Delta)$ be a channel machine, and let $s_F \in S$. Let $\mathcal{A}$ be the timed visibly one-counter net $\mathcal{A}$ in Fig. 2, and let $\mathcal{B}$ be the single-clock timed automaton such that $L(\mathcal{B}) = T\Sigma^+ \backslash L(\mathcal{C})$ as explained above. By Corollary 2, there exists some error-free computation of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$ if, and only if, $L(\mathcal{A}) \cap L(\mathcal{C}) \neq \emptyset$. The latter is equivalent to $L(\mathcal{A}) \not\subseteq L(\mathcal{B})$. Hence, the language inclusion problem is undecidable. □

4.5 Undecidability of the Model Checking Problem for MTL

The proof idea of Th. 1 can be reused to show the undecidability of the model checking problem for timed visibly one-counter nets and MTL over finite timed words. We briefly recall the syntax and semantics of MTL in the following.

Let Σ be a finite alphabet. The set of MTL formulae is built up from Σ by boolean connectives and a constraining version of the *until* modality:

$$\varphi ::= a \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \mathsf{U}_I \varphi_2$$

where $a \in \Sigma$ and $I \subseteq \mathbb{R}_{\geq 0}$ is an open, closed, or half-open interval with endpoints in $\mathbb{N} \cup \{\infty\}$.

We interpret MTL formulae in the *pointwise semantics*, *i.e.*, over finite timed words over Σ. Let $w = (a_1, t_1)(a_2, t_2) \dots (a_n, t_n)$ be a timed word, and let $i \in \{1, \dots, n\}$. We define the *satisfaction relation for MTL*, denoted by $\models$, inductively as follows:

$$\begin{aligned}
(w,i) \models a &\;\Leftrightarrow\; a_i = a \\
(w,i) \models \neg\varphi &\;\Leftrightarrow\; (w,i) \not\models \varphi, \\
(w,i) \models \varphi_1 \wedge \varphi_2 &\;\Leftrightarrow\; (w,i) \models \varphi_1 \text{ and } (w,i) \models \varphi_2, \\
(w,i) \models \varphi_1 \mathsf{U}_I \varphi_2 &\;\Leftrightarrow\; \exists j.\, i < j \leq |w| : (w,j) \models \varphi_2 \text{ and } t_j - t_i \in I, \\
&\qquad \text{and } \forall k.\, i < k < j : (w,k) \models \varphi_1.
\end{aligned}$$

We say that a timed word $w \in T\Sigma^+$ satisfies an MTL formula φ, written $w \models \varphi$, if $(w, 1) \models \varphi$.

Model Checking Problem for Timed Pushdown Automata and MTL
INPUT: A timed pushdown automaton $\mathcal{A}$, an MTL formula φ.

QUESTION: Does $w \models \varphi$ hold for every $w \in L(\mathcal{A})$?

Recall that this problem is decidable for the class of timed automata [30]. We can however prove that adding a visibly counter without zero test already makes the problem undecidable. Note that MTL only allows to express restrictions on time, and it does not allow for any restrictions on the values of the counters. In fact, it is known that as soon as we add to MTL the capability to express restrictions on the values of a counter that can be incremented and decremented, model checking is undecidable [32].

Theorem 4. *The model checking problem for timed visibly one-counter nets and MTL is undecidable, even if the timed visibly one-counter net does not use any clocks and is deterministic.*

Proof. We define an MTL formula φ such that $L(\varphi) = L(\mathcal{C})$. By Cor. 2, there exists some error-free computation of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$ if, and only if, $L(\mathcal{A}) \cap L(\varphi) \neq \emptyset$. The latter, however, is equivalent to saying that there exists some timed word $w \in L(\mathcal{A})$ such that $w \not\models \varphi$. Hence the model checking problem is undecidable. The definition of φ is straightforward. □

We would like to remark that the proof of Th. 4 shares some similarities with a recent [15] result on the undecidability of model checking one-counter machines (*i.e.*, one-counter automata without input alphabet) and Freeze LTL with one register ($\mathrm{LTL}_1^{\downarrow}$, for short). In [16], it is proved that $\mathrm{LTL}_1^{\downarrow}$ is capable to encode computations of *counter automata with incrementing errors*. Similar to the situation for MTL and channel machines, $\mathsf{LTL}_1^{\downarrow}$ is however not capable to encode *error-free* computations of counter automata. In [15], a one-counter machine is used to repair this incapability, resulting in the undecidability of the model checking problem. The one-counter machine in [15] does not use zero tests; however, we point out that in contrast to our visibly timed one-counter net the one-counter machine in [15] is *non-deterministic*. Indeed, model checking *deterministic* one-counter machines and $\mathrm{LTL}_1^{\downarrow}$ is decidable [15].

4.6 Energy Problems on Timed Automata with Discrete Weights

Next we will consider an interesting extension of *lower-bound energy problems on weighted timed automata*, introduced in [10], which gained attention in the last years, see, *eg.*, [11,31,9]. In lower-bound energy problems, one is interested whether in a given automaton with some weight variable whose value can be increased and decreased, there exists a successful run in which all accumulated weight values are never below zero. Similar problems have also been considered for untimed settings, *eg.*, [27,19,20,13].

A *timed automaton with discrete weights* (dWTA, for short) is syntactically the same as a timed one-counter net. In the semantical graph induced by a dWTA, however, we allow the value of the counter (or, the *weight variable*) to become negative. Hence the value of the weight variable does not influence the behaviour of the dWTA, because, different to timed one-counter nets, transitions

that result in negative values are not blocked. Formally, a *state* of a dWTA $\mathcal{A}$ is a triple (l, ν, u) where $l \in \mathcal{L}$, $\nu \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$, and $u \in \mathbb{Z}$. We use $\mathcal{G}^{\mathcal{A}}$ to denote the set of states of $\mathcal{A}$. A dWTA $\mathcal{A}$ induces a transition relation $\Rightarrow_{\mathcal{A}}$ on $(\mathcal{G}^{\mathcal{A}} \times \mathbb{R}_{\geq 0} \times \Sigma \times \mathcal{G}^{\mathcal{A}})$ as follows: $\langle (l, \nu, u), \delta, a, (l', \nu', u') \rangle \in \Rightarrow_{\mathcal{A}}$ if, and only if, there exists some edge $e = (l, a, \phi, \mathsf{op}, \lambda, l') \in E$ such that $(\nu + \delta) \models \phi$, $\nu' = (\nu + \delta)[\lambda := 0]$, and (i) if $\mathsf{op} = pop$, then $u' = u - 1$, and (ii) if $\mathsf{op} = push$, then $u' = u + 1$. *Runs*, *successful runs*, *associated timed words*, and the *language* $L(\mathcal{A})$ *accepted by* $\mathcal{A}$ are defined as for timed pushdown automata. We remark that due to the simple reason that the value of the weight variable does not influence the behaviour of dWTA, the model checking problem for dWTA and MTL is decidable.

Energy Problem for Timed Automata with Discrete Weights and MTL
INPUT: A dWTA $\mathcal{A}$, an MTL formula φ.
QUESTION: Does there exist some accepting run ρ of $\mathcal{A}$ such that the value of the weight variable is always non-negative, and the timed word w associated with ρ satisfies φ?

Theorem 5. *The energy problem for dWTA and MTL is undecidable, even if the dWTA uses no clock.*

Proof. For the proof, we reduce the model checking problem for timed one-counter nets and MTL to the energy problem. Let $\mathcal{A}$ be a timed one-counter net, and let φ be an MTL formula. Define $\mathcal{A}'$ to be the dWTA that is syntactically equal to $\mathcal{A}$. We prove that $(\mathcal{A}, \varphi)$ is a negative instance of the model checking problem if, and only if, $(\mathcal{A}', \neg\varphi)$ is a positive instance of the energy problem.

Assume that $(\mathcal{A}, \varphi)$ is a negative instance of the model checking problem. Hence there is some $w \in L(\mathcal{A})$ such that $w \models \neg\varphi$. Hence there is some successful run ρ of $\mathcal{A}$ such that ρ is associated with w, and by the semantics of $\mathcal{A}$, the value of the counter in ρ is always non-negative. Hence ρ is also an accepting run of $\mathcal{A}'$ on w such that the value of the weight variable is always non-negative and the associated timed word w satisfies $\neg\varphi$.

Assume that $(\mathcal{A}', \neg\varphi)$ is a positive instance of the energy problem. Hence there is some successful run ρ of $\mathcal{A}'$ such that the weight variable is always non-negative, and the timed word w associated with ρ satisfies $\neg\varphi$. Hence ρ is also a successful run of $\mathcal{A}$ on w. Hence $(\mathcal{A}, \varphi)$ is a negative instance of the model checking problem. □

5 Decidability Result

In this section, we prove that the language inclusion problem is decidable if $\mathcal{A}$ is a timed automaton, and $\mathcal{B}$ is a single-clock timed counter net. The proof is a generalization of the proof for timed automata in [29]. It is based on the theory of well-quasi-orders, and we start with defining some useful notions.

Let A, B be two sets, and let $\preceq$ be a binary relation on A. Then $\preceq$ is a *quasi-order* on A if $\preceq$ is reflexive and transitive. $\preceq$ is a *well-quasi-order* on A if

it is a quasi-order and for every infinite sequence $a_1, a_2, a_3, \ldots$ in A there exist indices $i < j$ such that $a_i \preceq a_j$. A standard example for a well-quasi-order is the component-wise ordering $\leq^k$ on the set $\mathbb{N}^k$ of vectors of k natural numbers (Dickson's Lemma, [17]).

Let $\preceq$ be a quasi-order on A, and let $\sqsubseteq$ be a quasi-order on B. We define the *product* of $\preceq$ and $\sqsubseteq$ on $(A \times B)$ by $(a, b) \leq (a', b')$, if and only if, $a \preceq a'$ and $b \sqsubseteq b'$. We define the *monotone domination order* $\preceq^*$ on A^* by $a_1 a_2 \ldots a_m \preceq^* a'_1 a'_2 \ldots a'_n$ if there exists a strictly increasing function $f : \{1, \ldots, m\} \to \{1, \ldots, n\}$ such that for all $i \in \{1, \ldots, m\}$ we have $a_i \preceq a'_{f(i)}$.

Lemma 4 (Higman's Lemma [22]).

1. *If $\preceq$ and $\sqsubseteq$ are well-quasi-orders on A and B, respectively, then the product of $\preceq$ and $\sqsubseteq$ is a well-quasi-order on $(A \times B)$.*
2. *If $\preceq$ is a well-quasi-order on A, then the monotone domination order $\preceq^*$ is a well-quasi-order on A^*.*

Proof of Theorem 2 For the case that $\mathcal{A}$ is a timed automaton and $\mathcal{B}$ is a single-clock timed automaton, the language inclusion problem $L(\mathcal{A}) \subseteq L(\mathcal{B})$ is decidable [29] with non-primitive recursive complexity [2]. The lower bound hence follows, and the decidability proof is an adaptation of the decidability proof in [29]. We present the idea, but do not give the details of the proofs, especially if they are similar to the proof in [29]. Let $\mathcal{A} = (\Sigma^{\mathcal{A}}, \Gamma^{\mathcal{A}}, \mathcal{L}^{\mathcal{A}}, \mathcal{L}_0^{\mathcal{A}}, \mathcal{L}_f^{\mathcal{A}}, \mathcal{X}, E^{\mathcal{A}})$ be a timed automaton with clock variables $\mathcal{X} = \{x_1, \ldots, x_n\}$, and let $\mathcal{B} = (\Sigma^{\mathcal{B}}, \mathfrak{n}, \mathcal{L}^{\mathcal{B}}, \mathcal{L}_0^{\mathcal{B}}, \mathcal{L}_f^{\mathcal{B}}, \{x\}, E^{\mathcal{B}})$ be a timed one-counter net of dimension $\mathfrak{n}$ with a single clock variable x. Without loss of generality, we may assume that $\Sigma^{\mathcal{A}} = \Sigma^{\mathcal{B}}$.

Note that a state (l, ν) of $\mathcal{A}$ is an element in $\mathcal{L}^{\mathcal{A}} \times (\mathbb{R}_{\geq 0})^{\mathcal{X}}$, and a state $(l, v, \boldsymbol{u})$ of $\mathcal{B}$ is an element in $\mathcal{L}^{\mathcal{B}} \times \mathbb{R}_{\geq 0} \times \mathbb{N}^{\mathfrak{n}}$. A *joint configuration of $\mathcal{A}$ and $\mathcal{B}$* is a pair (q, γ), where q is a state of $\mathcal{A}$, and γ is a set of states of $\mathcal{B}$. We use $\mathcal{G}^{\mathcal{AB}}$ to denote the set of all joint configurations of $\mathcal{A}$ and $\mathcal{B}$. We say that a joint configuration (q, γ) is *initial* if $q \in (\mathcal{L}_0^{\mathcal{A}} \times \{0\}^{\mathcal{X}})$ and $\gamma = \{(l_0^{\mathcal{B}}, 0, \mathbf{0})\}$. We say that a joint configuration (q, γ) is *bad* if $q = (l, \nu)$ for some $l \in \mathcal{L}_f^{\mathcal{A}}$, and for all states $(l', v', \boldsymbol{u'}) \in \gamma$ we have $l' \notin \mathcal{L}_f^{\mathcal{B}}$. The *joint behaviour of $\mathcal{A}$ and $\mathcal{B}$* is defined as follows: For a state (l, ν) of $\mathcal{A}$, $\delta \in \mathbb{R}_{\geq 0}$ and $a \in \Sigma$, we define $\mathsf{Succ}^{\mathcal{A}}((l, \nu), \delta, a) = \{(l', \nu') \mid \langle (l, \nu), \delta, a, (l', \nu') \rangle \in \Rightarrow_{\mathcal{A}}\}$. For a set γ of states of $\mathcal{B}$, we define $\mathsf{Succ}^{\mathcal{B}}(\gamma, \delta, a) = \{(l', v', \boldsymbol{u'}) \mid \exists (l, v, \boldsymbol{u}) \in \gamma . \langle (l, v, \boldsymbol{u}), \delta, a, (l', v', \boldsymbol{u'}) \rangle \in \Rightarrow_{\mathcal{B}}\}$. Note that $\mathsf{Succ}^{\mathcal{B}}(\gamma, \delta, a)$ is a set of states of $\mathcal{B}$, and it may be empty. Finally, we define the transition relation $\Rightarrow_{\mathcal{AB}}$ on $(\mathcal{G}^{\mathcal{AB}} \times \Sigma \times \mathcal{G}^{\mathcal{AB}})$ by $\langle (q, \gamma), a, (q', \gamma') \rangle \in \Rightarrow_{\mathcal{AB}}$ if there exists some $\delta \in \mathbb{R}_{\geq 0}$ such that $q' \in \mathsf{Succ}^{\mathcal{A}}(q, \delta, a)$ and $\gamma' = \mathsf{Succ}^{\mathcal{B}}(\gamma, \delta, a)$.

Next, we encode joint configurations of $\mathcal{A}$ and $\mathcal{B}$ by finite untimed words over the alphabet Λ of finite subsets of $(\mathcal{L}^{\mathcal{A}} \times \mathcal{X} \times \mathsf{reg} \times \mathbf{0}) \cup \{\mathcal{L}^{\mathcal{B}} \times \{x\} \times \mathsf{reg} \times \mathbb{N}^{\mathfrak{n}})$. Here, $\mathbf{0}$ is the vector of dimension $\mathfrak{n}$ containing only 0, and $\mathsf{reg} := \{0, 1, \ldots, \mathsf{cmax}\} \cup \{\top\}$, where $\top$ is a symbol representing all values greater than cmax, and cmax is an integer greater than the maximal constant occurring in clock constraints in both $\mathcal{A}$ and $\mathcal{B}$. Let $C = ((l, \nu), \{(l_1, v_1, \boldsymbol{u_1}), \ldots, (l_m, v_m, \boldsymbol{u_m})\})$ be a joint

configuration. To simplify the definition, we write C as a set

$$\{(l, x_1, \nu(x_1), \mathbf{0}), \ldots, (l, x_n, \nu(x_n), \mathbf{0}), (l_1, x, v_1, \boldsymbol{u_1}), \ldots, (l_m, x, v_m, \boldsymbol{u_m})\}.$$

Split C into a sequence of subsets $C_0, C_1, \ldots, C_\rho, C_\top$, such that $C_\top = \{(k, y, \eta, \boldsymbol{\mu}) \in C \mid \eta > \mathsf{cmax}\}$, and if $i, j \neq \top$, then for all $(k, y, \eta, \boldsymbol{\mu}) \in C_i$, $(k', y', \eta', \boldsymbol{\mu'}) \in C_j$, we have $\mathsf{frac}(\eta) \leq \mathsf{frac}(\eta')$ if, and only if, $i \leq j$. In this way, $(k, y, \eta, \boldsymbol{\mu})$ and $(k', y', \eta', \boldsymbol{\mu'})$ are in the same subset C_i if, and only if, η and η' are both smaller than or equal to cmax and have the same fractional part. In addition, we assume that the fractional parts of η of elements in C_0 are zero, and that $C_i \neq \emptyset$ for all $i \in \{1, \ldots, \rho\}$. We define the encoding $\mathsf{enc}(C)$ of C to be the finite word $\mathsf{reg}(C_0)\mathsf{reg}(C_1)\ldots\mathsf{reg}(C_\rho)\mathsf{reg}(C_\top)$, where $\mathsf{reg}(C_i) = \{(k, y, \mathsf{reg}(\eta), \boldsymbol{\mu}) \mid (k, y, \eta, \boldsymbol{\mu}) \in C_i\}$ with $\mathsf{reg}(\eta) = \mathsf{int}(\eta)$ if $\eta \leq \mathsf{cmax}$, and $\mathsf{reg}(\eta) = \top$ otherwise.

We define a transition relation $\rightarrow$ over encodings of joint configurations and Σ as follows: $\langle w, a, w'\rangle \in \rightarrow$ if there exists $C \in \mathsf{enc}^{-1}(w)$ and $C' \in \mathsf{enc}^{-1}(w')$ such that $\langle C, a, C'\rangle \in \Rightarrow_{\mathcal{AB}}$.

Lemma 5. *The equivalence relation $\sim$ defined by $C \sim C :\Leftrightarrow \mathsf{enc}(C) = \mathsf{enc}(C')$ is a time-abstract bisimulation over joint configurations.*

Like in [29], we can now define a relation $\preceq$ over the set of encodings of joint configurations. The letters of our alphabet are naturally ordered by set inclusion, so that we can define $\preceq$ to be the monotone domination order induced by set inclusion. In [29], using Higman's Lemma and the fact that the alphabet is finite, it is proved that $\preceq$ is a well-quasi-order. However, note that in contrast to [29], our alphabet is infinite. We thus define a new relation $\sqsubseteq$ over the set of encodings, which is basically the product of $\preceq$ and the usual $\leq^{\mathtt{n}}$ relation over $\mathbb{N}^{\mathtt{n}}$. Since both $\preceq$ and $\leq^{\mathtt{n}}$ are well-quasi-orders, by Higman's Lemma $\sqsubseteq$ is a well-quasi-order, too.

Formally, we define $\sqsubseteq$ as follows. Let $w = \alpha_1\alpha_2\ldots\alpha_m$ and let $w = \beta_1\beta_2\ldots\beta_p$, where $\alpha_i = \{(k^i_1, y^i_1, \eta^i_1, \boldsymbol{\mu^i_1}), \ldots, (k^i_{i_\rho}, y^i_{i_\rho}, \eta^i_{i_\rho}, \boldsymbol{\mu^i_{\rho_i}})\}$, and $\beta_i = \{(l^i_1, z^i_1, \zeta^i_1, \boldsymbol{\omega^i_1}), \ldots, (l^i_{i_\rho}, z^i_{i_\rho}, \zeta^i_{i_\rho}, \boldsymbol{\omega^i_{\rho_i}})\}$. We define $w \sqsubseteq w'$ if, and only if, there exists an injective function $f : \{1, \ldots, m\} \rightarrow \{1, \ldots, p\}$ such that for every $i \in \{1, \ldots, m\}$, there is an injective function $f_i : \{1, \ldots, \rho_i\} \rightarrow \{1, \ldots, \omega_{\rho_{f(i)}}\}$ that satisfies $k^i_j = l^{f(i)}_{f_i(j)}$, $y^i_j = z^{f(i)}_{f_i(j)}$, $\eta^i_j = \zeta^{f(i)}_{f_i(j)}$, and $\mu^i_j \leq^{\mathtt{n}} \omega^{f(i)}_{f_i(j)}$ for every $j \in \{1, \ldots, \rho_i\}$. Next we prove that $\rightarrow$ is downward-compatible with respect to $\sqsubseteq$.

Lemma 6. *If $w_1 \sqsubseteq w_2$ and $\langle w_2, a, w_2'\rangle \in \rightarrow$, then there exists $w_1' \sqsubseteq w_2'$ such that $\langle w_1, a, w_1'\rangle \in \rightarrow$.*

Finally, we describe the algorithm to decide $L(\mathcal{A}) \subseteq L(\mathcal{B})$. Like in [29], we solve the language inclusion problem by solving the following reachability problem: in the implicit graph of the encoding of joint configurations and the transition relation $\rightarrow$, is there a path from the encoding of the initial joint configuration to the encoding of a bad joint configuration? Note that we have $L(\mathcal{A}) \subseteq L(\mathcal{B})$ if, and only if, there is no such path. For solving the reachability problem, we compute the unfolding of the graph, starting the computation with the encoding

of the initial joint configuration. If for the current node labelled by w, there is along the branch already a node labelled with w' and $w' \sqsubseteq w$, then by Lemma 6 we can prune the tree after the current node: Assume that from w we can reach a word w_1 that represents a bad configuration, then by Lemma 6 we can reach a word w_1' from w' such that $w_1' \sqsubseteq w_1$, and hence, w_1' is representing a bad configuration, too. By the facts that the unfolding is finitely branching, $\sqsubseteq$ is a well-quasi-order and by König's Lemma, we know that the computation will finally terminate. □

Note that Lemma 6 does not hold if the counters in $\mathcal{B}$ can be tested for zero. Indeed, the universality problem of (even untimed) one-counter automata is undecidable [21,24]. In the next section, we prove that this is also the case for timed visibly pushdown one-counter automata. This gives us the precise decidability border for the universality problem.

6 The Universality Problem for Visibly One-Counter Automata

We prove that allowing the counter in a single-clock timed visibly one-counter net to be tested for zero, results in the undecidability of the universality problem. The undecidability of the universality problem for the more general class of single-clock visibly pushdown automata was already stated in Theorem 3 in [18]. The proof is a reduction of the halting problem for two-counter machines. Given a two-counter machine $\mathcal{M}$, one can define a timed language $L(\mathcal{M})$ that consists of all timed words encoding a halting computation of $\mathcal{M}$. Then a timed visibly pushdown automaton $\mathcal{A}$ is defined that recognizes the complement of $L(\mathcal{M})$. Altogether, $L(\mathcal{A}) = T\Sigma^+$ if, and only if, $\mathcal{M}$ does not have a halting computation.

The definition of $L(\mathcal{M})$ is similar to the definition of $L(\mathcal{C})$ in the proof of Th. 1. Recall that in the definition of $L(\mathcal{C})$ we did not include a condition that requires every symbol to have a matching symbol two time units *before*, and, as we mentioned, this is the reason for $L(\mathcal{C})$ to contain timed words encoding *faulty* computations of $\mathcal{C}$. However, in the definition of $L(\mathcal{M})$ in [18], such a "backward-looking" condition is used. In the proof in [18], it is unfortunately not clear how the single-clock timed visibly pushdown automaton $\mathcal{A}$ can detect violations of this condition[2].

Here, we give a complete proof for the subclass of timed visibly one-counter automata. Like the proof of Th. 1, the proof is a reduction of the control state reachability problem for channel machines. We however remark that one can similarly use a reduction of the halting problem for two-counter machines.

[2] More detailed, it is not clear how to construct single-clock timed automata $N_{\neg f_r \leftarrow f_c}$ and $N_{\neg g_r \leftarrow g_c}$ mentioned on p. 10 in [18]. Recall that in the proof for undecidability of the universality problem for timed automata with two or more clocks, it is exactly this backward-looking condition that requires *two* clocks [3].

Proof of Theorem 3 The proof reuses ideas of the proof of Th. 1. Let $\mathcal{C} = (S, s_I, M, \Delta)$ be a channel machine, and let $s_F \in S$. Define Σ in the same way as in the proof of Th. 1. For every $n \in \mathbb{N}$, we define a timed language $L_{\mathsf{ef}}(\mathcal{C}, n)$ that consists of all timed words over Σ that encode *error-free* computations of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$. Formally, $L_{\mathsf{ef}}(\mathcal{C}, n)$ is defined using the same conditions as the ones for $L(\mathcal{C}, n)$ in the proof of Th. 1. However, we add the following condition to exclude encodings of faulty computations:

14. For every infix of w of the form

$$(s, \delta)(\sigma_1, \delta_1) \dots (\sigma_k, \delta_k)(l, \delta + 1)(s', \delta + 2)(\sigma'_1, \delta'_1) \dots (\sigma'_{k'}, \delta'_{k'})(l, \delta + 3)$$

with $s, s' \in S$, $l \in L$, $l' \in L \cup \{\star\}$, $\delta \in \mathbb{R}_{\geq 0}$, $\delta < \delta_1 < \dots < \delta_k < \delta + 1$, $\delta + 2 < \delta'_1 < \dots < \delta'_{k'} < \delta + 3$, there exists a strictly increasing function $f : \{1, \dots, k'\} \to \{1, \dots, k\}$ such that $\delta_{f(i)} = \delta'_i - 2$.

We define $L_{\mathsf{ef}}(\mathcal{C}) = \bigcup_{n \geq 1} L_{\mathsf{ef}}(\mathcal{C}, n)$.

Lemma 7. *There exists some error-free computation of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$, if, and only if, $L_{\mathsf{ef}}(\mathcal{C}) \neq \emptyset$.*

Proof. The direction from right to left is straightforward. For the other direction assume γ is an error-free computation of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$. Put $n = \max(\gamma)$. By Lemma 1, there exists $w \in L(\mathcal{C}, n)$ such that $\max(w) = n$. A close inspection of conditions 11, 12, 13 and their primed versions reveals that we can also ensure that condition 14 is satisfied. Hence $w \in L_{\mathsf{ef}}(\mathcal{C})$. □

Next, we define a timed visibly one-counter automaton with a single clock such that $L(\mathcal{A}) = T\Sigma^+ \backslash L_{\mathsf{ef}}(\mathcal{C})$. Hence, by the preceding lemma, $L(\mathcal{A}) \neq T\Sigma^+$ if, and only if, there exists some error-free computation of $\mathcal{C}$ from (s_I, ε) to (s_F, x) for some $x \in M^*$.

$\mathcal{A}$ is the union of several single-clock timed automata and one timed visibly one-counter automaton with no clock. We already know from the proof of Th. 1, that violations of conditions 1 to 13 can be detected by single-clock timed automata. For detecting violations of condition 14, we use the visibly one-counter automaton shown in Fig. 4. The automaton non-deterministically guesses the maximum number n of occurrences of the symbol $+$. When leaving l_1, the value of the counter is $n + 1$. The final location l_4, however, can only be reached while reading $-$ or $\star$ if the value of the counter is zero. This means that there must be some symbol for which there is no matching symbol one time unit before. □

7 Conclusion and Open Problems

The main conclusion of this paper is that even very weak extensions of timed automata with counters are impossible to be verified by automated algorithms. On the other hand side, we may use single-clock timed counter nets as specifications to verify timed automata. This increases so far known possibilities for the verification of timed automata: For instance, the timed language $L = \{(a^m b^n, \bar{\tau}) \mid$

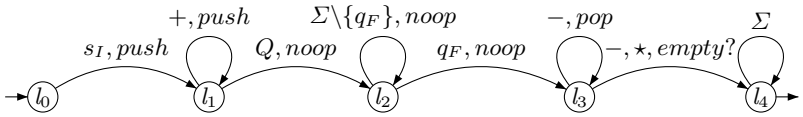

Fig. 4. The timed visibly one-counter automaton for recognizing timed words violating condition 14 of $L_{\mathsf{ef}}(\mathcal{C})$.

$m \geq n\}$ can be accepted by a timed one-counter net without any clocks, but not by a timed automaton.

An interesting problem is to figure out a (decidable) extension of LTL that is capable to express properties referring to both time and discrete data structures.

We remark that all our results hold for automata defined over *finite* timed words. We cannot expect the decidability of, *eg.*, the universalitiy problem for single-clock timed counter nets over infinite timed words, as the same problem is already undecidable for the subclass of single-clock timed automata [2].

Acknowledgements

I would like to thank Michael Emmi and Rupak Majumdar for helpful discussions on their work on timed pushdown automata. I further would like to thank James Worrell very much for pointing me to MTL's capability to encode faulty computations of channel machines.

References

1. Parosh Aziz Abdulla and Karlis Cerans. Simulation is decidable for one-counter nets (extended abstract). In Davide Sangiorgi and Robert de Simone, editors, *CONCUR*, volume 1466 of *LNCS*, pages 253–268. Springer, 1998.
2. Parosh Aziz Abdulla, Johann Deneux, Joël Ouaknine, Karin Quaas, and James Worrell. Universality analysis for one-clock timed automata. *Fundam. Inform.*, 89(4):419–450, 2008.
3. Sara Adams, Joël Ouaknine, and James Worrell. Undecidability of universality for timed automata with minimal resources. In Jean-François Raskin and P. S. Thiagarajan, editors, *FORMATS*, volume 4763 of *LNCS*, pages 25–37. Springer, 2007.
4. Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
5. Rajeev Alur and P. Madhusudan. Decision problems for timed automata: A survey. In Marco Bernardo and Flavio Corradini, editors, *SFM*, volume 3185 of *LNCS*, pages 1–24. Springer, 2004.
6. Rajeev Alur and P. Madhusudan. Visibly pushdown languages. In László Babai, editor, *STOC*, pages 202–211. ACM, 2004.

7. Ahmed Bouajjani, Rachid Echahed, and Riadh Robbana. On the automatic verification of systems with continuous variables and unbounded discrete data structures. In Panos J. Antsaklis, Wolf Kohn, Anil Nerode, and Shankar Sastry, editors, *Hybrid Systems*, volume 999 of *LNCS*, pages 64–85. Springer, 1994.
8. Florent Bouchy, Alain Finkel, and Arnaud Sangnier. Reachability in timed counter systems. *Electr. Notes Theor. Comput. Sci.*, 239:167–178, 2009.
9. Patricia Bouyer, Uli Fahrenberg, Kim G. Larsen, and Nicolas Markey. Timed automata with observers under energy constraints. In Karl Henrik Johansson and Wang Yi, editors, *HSCC*, pages 61–70. ACM, 2010.
10. Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jirí Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *FORMATS*, volume 5215 of *LNCS*, pages 33–47. Springer, 2008.
11. Patricia Bouyer, Kim G. Larsen, and Nicolas Markey. Lower-bound-constrained runs in weighted timed automata. *Perform. Eval.*, 73:91–109, 2014.
12. Daniel Brand and Pitro Zafiropulo. On communicating finite-state machines. *J. ACM*, 30(2):323–342, April 1983.
13. T. Brázdil, P. Jancar, and A. Kucera. Reachability games on extended vector addition systems with states. In S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide, and P. G. Spirakis, editors, *ICALP (2)*, volume 6199 of *LNCS*, pages 478–489. Springer, 2010.
14. Zhe Dang. Pushdown timed automata: a binary reachability characterization and safety verification. *Theor. Comput. Sci.*, 302(1-3):93–121, 2003.
15. S. Demri, R. Lazić, and A. Sangnier. Model checking Freeze LTL over one-counter automata. In R. M. Amadio, editor, *FoSSaCS*, volume 4962 of *LNCS*, pages 490–504. Springer, 2008.
16. Stéphane Demri and Ranko Lazić. LTL with the freeze quantifier and register automata. *ACM Trans. Comput. Log.*, 10(3), 2009.
17. Leonard Eugene Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *Amer. J. Math.*, 35:413–422, 1913.
18. Michael Emmi and Rupak Majumdar. Decision problems for the verification of real-time software. In J. P. Hespanha and A. Tiwari, editors, *HSCC*, volume 3927 of *LNCS*, pages 200–211. Springer, 2006.
19. Zoltán Ésik, Uli Fahrenberg, Axel Legay, and Karin Quaas. Kleene algebras and semimodules for energy problems. In Dang Van Hung and Mizuhito Ogawa, editors, *ATVA*, volume 8172 of *LNCS*, pages 102–117. Springer, 2013.
20. Uli Fahrenberg, Line Juhl, Kim G. Larsen, and Jirí Srba. Energy games in multiweighted automata. In Antonio Cerone and Pekka Pihlajasaari, editors, *ICTAC*, volume 6916 of *LNCS*, pages 95–115. Springer, 2011.
21. Sheila A. Greibach. An infinite hierarchy of context-free languages. *J. ACM*, 16(1):91–106, 1969.
22. Graham Higman. Ordering by divisibility in abstract algebras. *Proceedings of the London Mathematical Society*, 2:236–366, 1952.
23. Piotr Hofman, Slawomir Lasota, Richard Mayr, and Patrick Totzke. Simulation over one-counter nets is PSPACE-complete. In Anil Seth and Nisheeth K. Vishnoi, editors, *FSTTCS*, volume 24 of *LIPIcs*, pages 515–526. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
24. Oscar H. Ibarra. Restricted one-counter machines with undecidable universe problems. *Mathematical Systems Theory*, 13:181–186, 1979.
25. Petr Jancar, Javier Esparza, and Faron Moller. Petri nets and regular processes. *J. Comput. Syst. Sci.*, 59(3):476–503, 1999.

26. Petr Jancar, Antonín Kucera, Faron Moller, and Zdenek Sawa. DP lower bounds for equivalence-checking and model-checking of one-counter automata. *Inf. Comput.*, 188(1):1–19, 2004.
27. Line Juhl, Kim Guldstrand Larsen, and Jean-François Raskin. Optimal bounds for multiweighted and parametrised energy games. In Zhiming Liu, Jim Woodcock, and Huibiao Zhu, editors, *Theories of Programming and Formal Methods*, volume 8051 of *LNCS*, pages 244–255. Springer, 2013.
28. Information Sciences Institute of the University of Southern California. *Transmission Control Protocoll* (DARPA Internet Program Protocol Specification), 1981. http://www.faqs.org/rfcs/rfc793.html.
29. Joël Ouaknine and James Worrell. On the language inclusion problem for timed automata: Closing a decidability gap. In *LICS*, pages 54–63. IEEE Computer Society, 2004.
30. Joël Ouaknine and James Worrell. On the decidability of metric temporal logic. In *LICS*, pages 188–197. IEEE Computer Society, 2005.
31. Karin Quaas. On the interval-bound problem for weighted timed automata. In Adrian Horia Dediu, Shunsuke Inenaga, and Carlos Martín-Vide, editors, *LATA*, volume 6638 of *LNCS*, pages 452–464. Springer, 2011.
32. Karin Quaas. Model checking metric temporal logic over automata with one counter. In Adrian Horia Dediu, Carlos Martín-Vide, and Bianca Truthe, editors, *LATA*, volume 7810 of *LNCS*, pages 468–479. Springer, 2013.